

**BLACKBURN
EXHIBIT I**

Determining User Location For Context Aware Computing Through the Use of a Wireless LAN Infrastructure

Jason Small
masterz@andrew.cmu.edu

Asim Smailagic
asim@cs.cmu.edu

Daniel P. Siewiorek
dps@cs.cmu.edu

*Institute for Complex Engineered Systems
Carnegie Mellon University, Pittsburgh, PA 15213*

Abstract

Determining the location of a mobile computer user has several applications. We have developed and demonstrated a person locator using the wireless network infrastructure at Carnegie Mellon University to accomplish this task. A user's location can be determined both indoors and outdoors while on campus, and at a higher resolution than the Global Positioning System. The system requires a minimum of extra hardware but results in a high degree of accuracy.

Keywords: User Location, Context Aware Computing, Wireless LAN

1. Introduction

Increased user mobility suggests that applications should adapt themselves based on knowledge of location. Location is an example of context, that is, information about people or devices that can be used to modify the way a system provides its services to the user community [1]. The use of location information to infer user intent based on location and previous user actions can enrich the capabilities of mobile users.

Currently, the most common method of determining a person's position is by using the Global Positioning System (GPS). GPS is used for applications ranging from aircraft and ship navigation to hikers and automobiles. The recent "undegradation" of the signal will allow for increased accuracy and use [2]. Unfortunately, GPS does not work indoors. For this reason, other methods must be used.

We have developed a method using a newly built wireless network infrastructure at Carnegie Mellon University by which a user's location can be determined both indoors and outdoors while on campus, and at a higher resolution than GPS. To the best of our knowledge, our method of a table - based lookup for triangulation of a user's location

using a wireless LAN has not been attempted before.

1.1. Context Aware Applications

Physical location is a particularly valuable attribute in context-aware computing. The machine which assists us most intelligently is the one which understands our real-world context. This context is quite rich, consisting of attributes such as our physical location, our state of mind, our personal history, our present company, and an uncountable number of other features. While making these decisions, the assistant could choose not to disturb the user at inopportune moments. Ultimately, the assistant could become the most specialized concierge, handling any number of tasks from personal scheduling to matchmaking to tour guiding [3].

Many applications for mobile wearable computers can be developed. Critical to the development of these services is location awareness, or the notion that the system knows the location of each user.

Our research in this area proceeds in two parts. First, we are developing an API that allows mobile clients to obtain and use location information in a technology independent manner. Second, we explore two specific approaches to location sensing based on the Wavelan 802.11 wireless network at Carnegie Mellon. The first approach consists of discovering the active access point for a mobile client and mapping that information onto a two-dimensional campus map. An access point covers a sphere of approximately 75 feet in diameter. The second approach improves resolution by triangulation based on measured signal strength from several nearby nodes.

1.2. Wireless Infrastructure

As part of Carnegie Mellon's Wireless Andrew initiative, all academic buildings on campus have been equipped

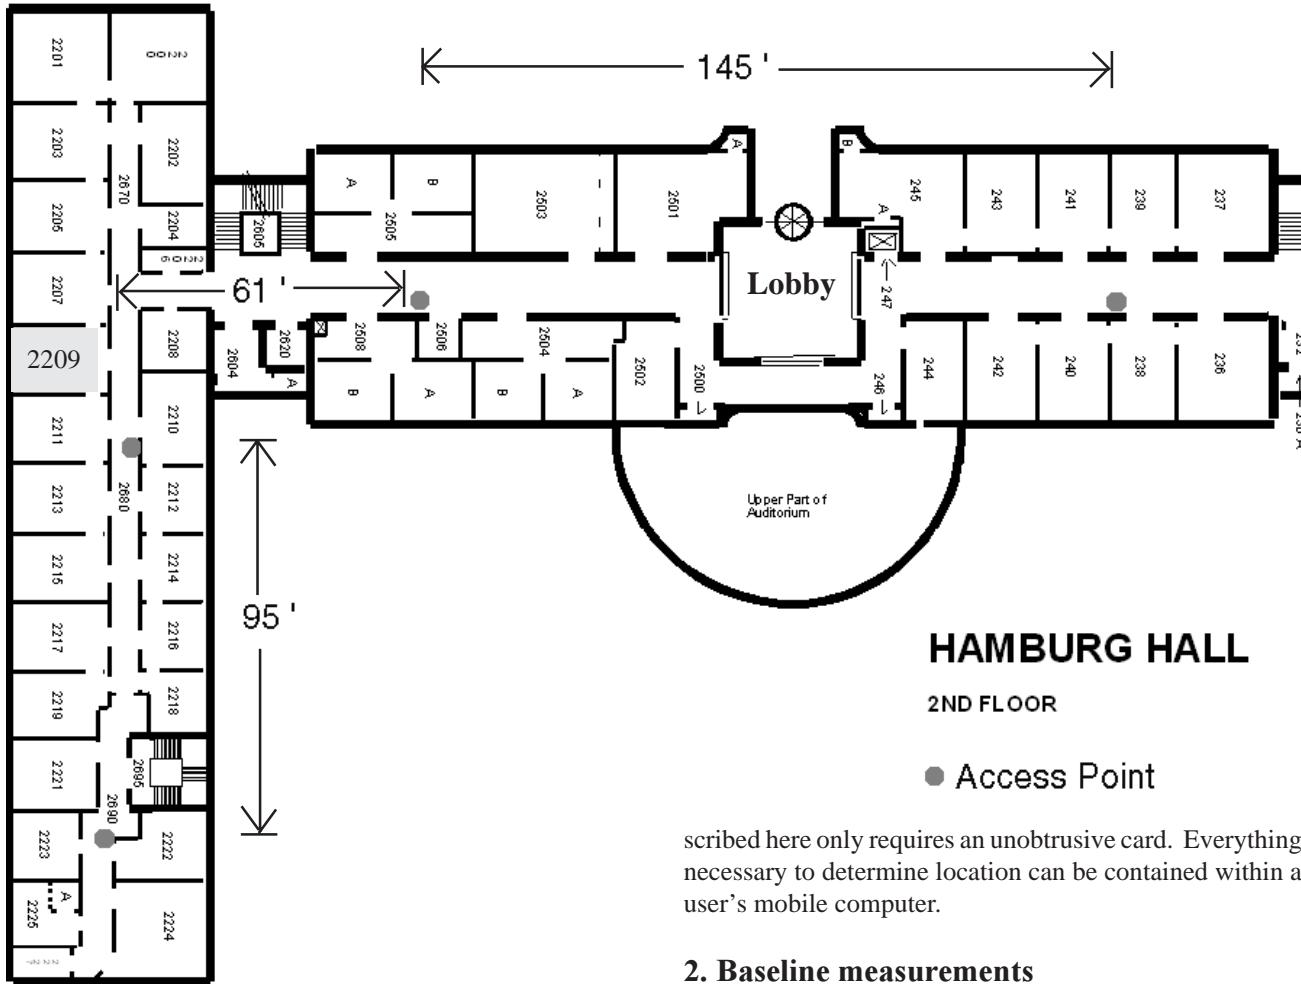


Figure 1. Layout of test locations

with wireless network connectivity. As of May 2000, there were approximately 400 wireless access points at Carnegie Mellon University, covering an area of 103 acres. Through these access points, members of the campus community can connect to the network anywhere on campus including buildings and outdoor spaces. Andrew network services can be accessed by a variety of wearable and handheld platforms. The first set of services explicitly for wireless handheld devices are under development by CMU's Computing Services and are being deployed. These services include electronic mail, scheduling and notification. These services will be extended with traffic information, and some other time sensitive information.

The wireless infrastructure is an enabler for next generation applications. Previously described methods of determining a user's location require external hardware such as a camera to be mounted on a user [4], or light meters, accelerometers and magnetometers [5]. The method de-

scribed here only requires an unobtrusive card. Everything necessary to determine location can be contained within a user's mobile computer.

2. Baseline measurements

Baseline measurements were taken to provide insight into the methods necessary to determine location. The layout of the testing area is depicted in Figure 1. Stationary measurements were taken in room 2209 on the left. Mobile measurements, were taken in the long hallway at the left. Other mobile measurements were taken in the hall to the left of the lobby.

2.1. Stationary measurements

The first objective was to determine whether signal strength was consistent for a specific location over time. If this was not the case, the reported location measurements would be inconsistent at different time. One way to overcome this obstacle would be to set up a system similar to differential GPS in which a stationary station provides correction data to handheld units.

Samples were taken over periods of five hours, 20 hours and a month at five - second intervals, Figure 2. The re-

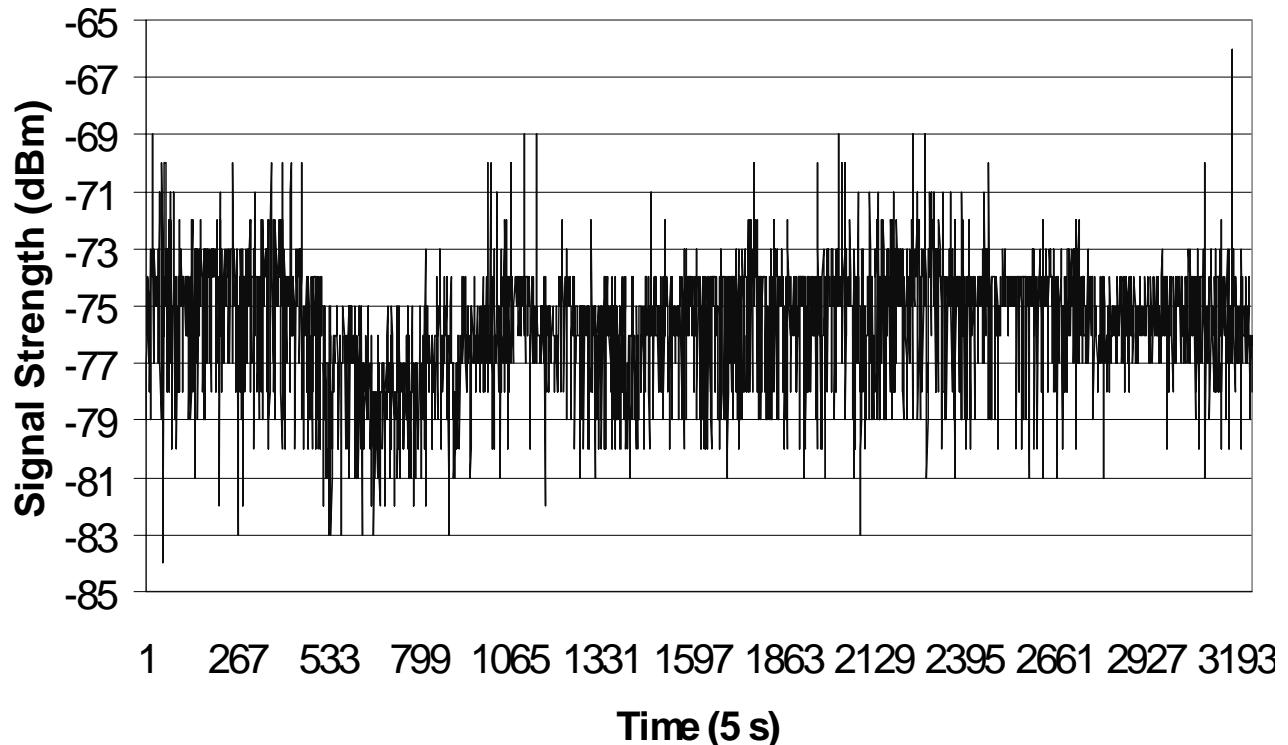


Figure 2. Signal strength measurements from a stationary position over time

sults show that the data is fairly consistent with a standard deviation of 2.13 dBm (decibel milliwatt), Figure 3. It was also determined that environmental condition such as opening an office door can create a change of up to 10 dBm (this was expected). What was not expected was the effect other handheld devices had on signal strength, having an effect of up to 5 dBm. This is what caused the dip during samples 524-984 in Figure 2.

A variation of +/- 5 dBm was measured for multiple samples at a single location, Figure 3. If we average 10

samples, this variation can be reduced to +/- 2.5 dBm. Averaging brings the standard deviation down to .939 dBm when within 20 feet of the access point and 1.146 dBm when 70 feet away. This puts 68.6% of the samples within +/- .939 dBm and 95.44% within +/- 2.292 dBm of the “actual” value.

2.2. Mobile measurements

The next objective was to determine to what degree signal strength changed with distance from the access point.

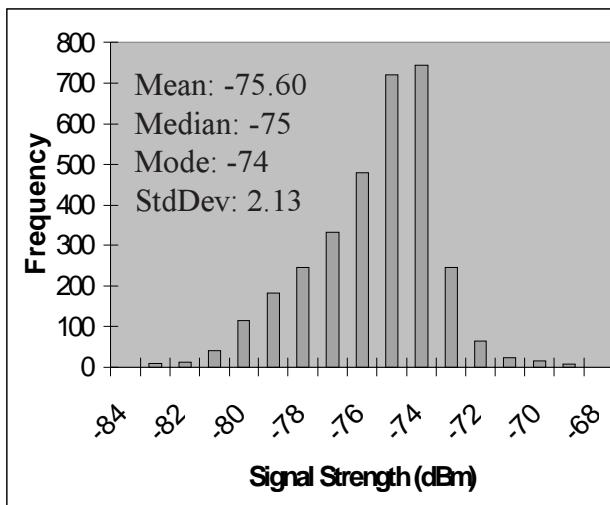


Figure 3. Distribution of signal strength

Accuracy	dBm	Feet
68.6 %	+/- .939	+/- 5
95.4 %	+/- 1.146	+/- 10
99.9 %	+/- 2.817	+/- 15

Table 1. Accuracy of location measurements

A quantum had to be measured to determine how much distance 1 dBm represents. This determines approximately what the maximum accuracy can be. The distance needed to result in a change of 1 dBm was determined to be 5 ft. while close to the access point with change in strength per foot decreasing further from the access point. This turns out not to be a large problem because by the time we are in a region where the strength is not changing much, we are already in range of another access point. Table 1 summarizes the accuracy of our results. Due to the fact that greater than 99.9% of our measurements are within 3 dBm of the actual value, we can be very confident that our readings are within +/- 15 feet of our actual position.

Standard deviation is most useful when dealing with a normal distribution. A normal distribution is characterised by data having the same mean, median and mode. The data in Figure 2 has a mean, median and mode of -75.6, -75 and -74 respectively, showing that the data follows a normal distribution. Two measures of variation are standard deviation and range. Figure 4, which was generated on multiple measurements at various locations shows that there is a linear relationship between range and standard deviation, indicating that larger ranges in values are not due to singular outliers, but to a general spread of the data.

General measurements were taken to see how distance correlated to signal strength. This is illustrated in Figure 5. Ten measurements were taken every five feet down a hallway, over a distance ranging from 10 to 135 feet. During this test, two access points were in range, with one situated at 40 feet down the hallway and the other at 135. These positions can be seen in the graph, as the signal strength peaks at these positions. Polynomial trend lines have been

added to make it easier to read the graph.

By sampling multiple times and averaging the values, the correlation (covariance divided by products of standard deviations between series) of data points from two data sets taken from various identical locations at different times was raised to .983, compared to .916 when the samples were not averaged. This allows us to discount stray readings and increases accuracy. If necessary, the accuracy can be increased even more by taking more samples and discarding the outliers.

It should be noted that the further away from the access point, the less accurate the readings. For this reason, it may be advantageous to place a higher degree of confidence in stronger signals and weight them accordingly. It should also be noted that while signal strength obtained closer to the access points are more accurate, the change in signal strength is greater per foot when closer to the access point. We are currently studying the effect of weighting certain readings to determine if accuracy can be improved.

One procedure, which has been determined to increase accuracy, is to accept only those results for which a sufficient number of samples have been taken and to ignore any results below a certain threshold. These usually turn out to be the same since there are fewer values returned for those access points that have low signal strength.

Measurements were also taken at different times of day. Figure 6 illustrates these measurements and shows that there is an insignificant difference between measurements taken at different times of day. It should be noted that while there were many more people in the vicinity of the measuring

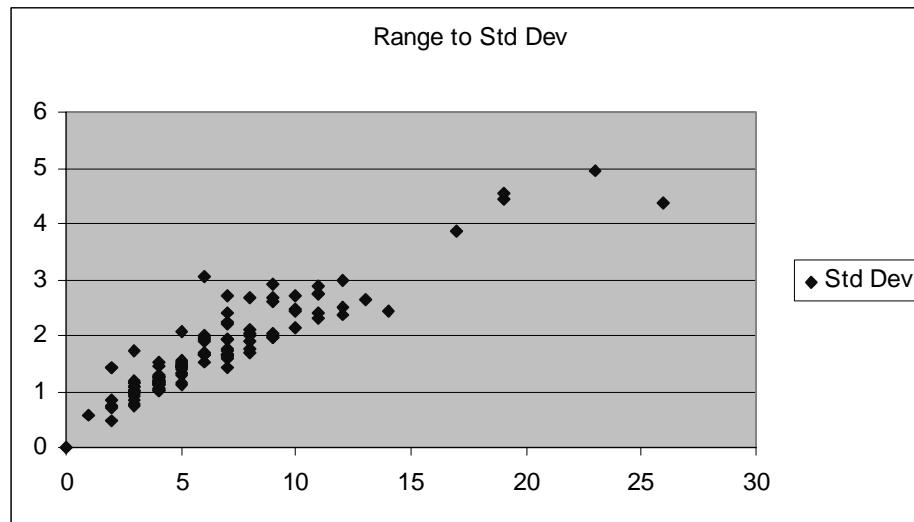


Figure 4. Relation of Range of Values to Standard Deviation of Values

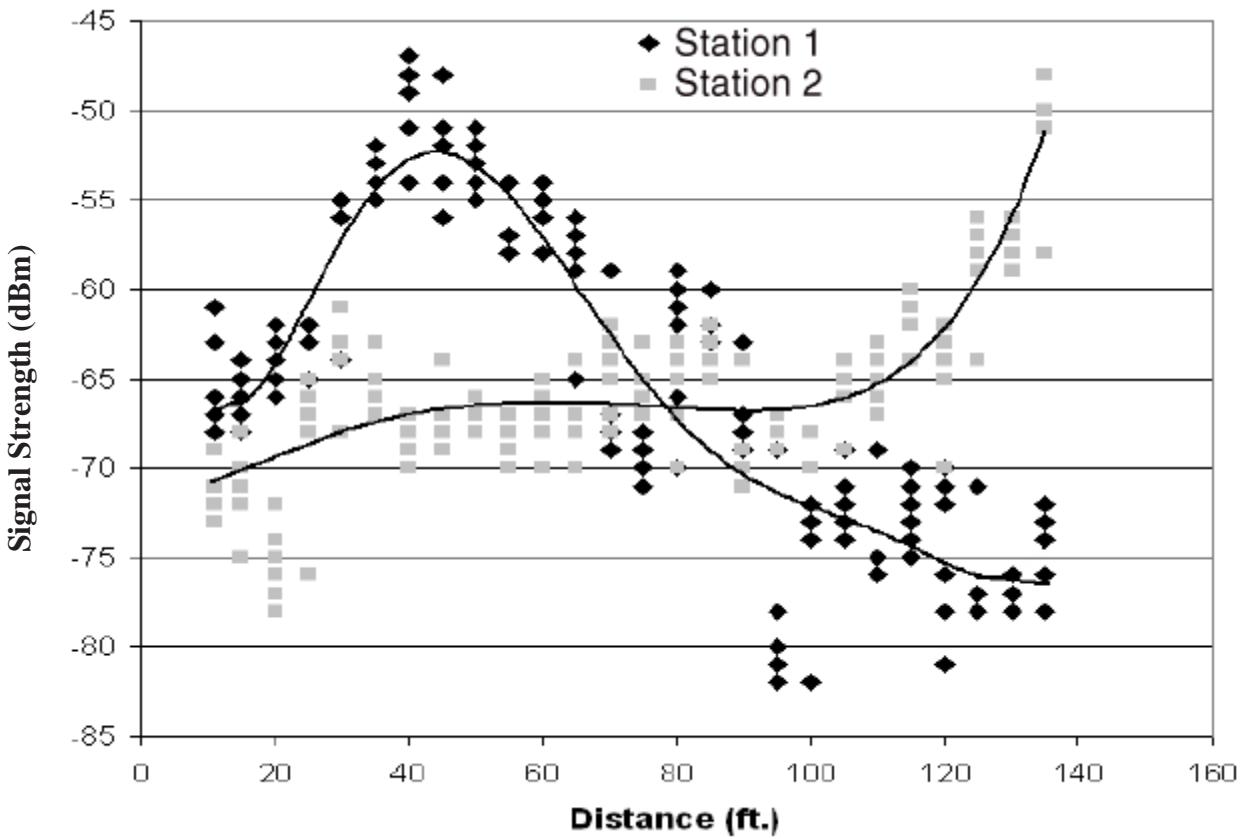


Figure 5. Relation of signal strength to distance

computer during the day, their presence did not affect the signal strength.

3. Methodology

Triangulation based on signal strength is used to determine the location of a user. For every location, there is a unique reading of signal strengths gathered from a group of access points. A minimum of one access point is necessary to determine location, with the addition of extra access points increasing accuracy.

Figure 7 shows the use of two access points to determine location. Each signal strength reading places the user within a ring around the access point. Using fewer than three access points can potentially give an ambiguous answer as seen in Figure 7. Readings place the user within the grey rings for each access points. It can be seen that these rings overlap in two areas. The addition of a third access point would allow for an unambiguous reading for the user's location.

Four methods have been explored to process this data: a table-based method delivering high resolution, a table-based

method delivering low resolution, a model-based method, and a method using a neural network.

3.1. Table - based: High Resolution

For table - based determination, samples must be taken in the area for which users will have the need to know their location. For this training, the user's location is manually input into the computer and about 17 samples are taken and averaged. A table is generated, allowing us to see what signal levels to expect at different locations. This only has to be done once and can be saved for use in later sessions and on other platforms.

During use, measured values are compared to those in the table and differences are computed. The entry with the smallest difference is taken to be the current position.

Measurements were taken at eight discrete locations along the hallway to the left of the lobby in Figure 1. The table was first populated with data from each location. After the table was populated, the computer was taken back to each previous location and using this method, reported the computed location. The results are summarized in Table 2.

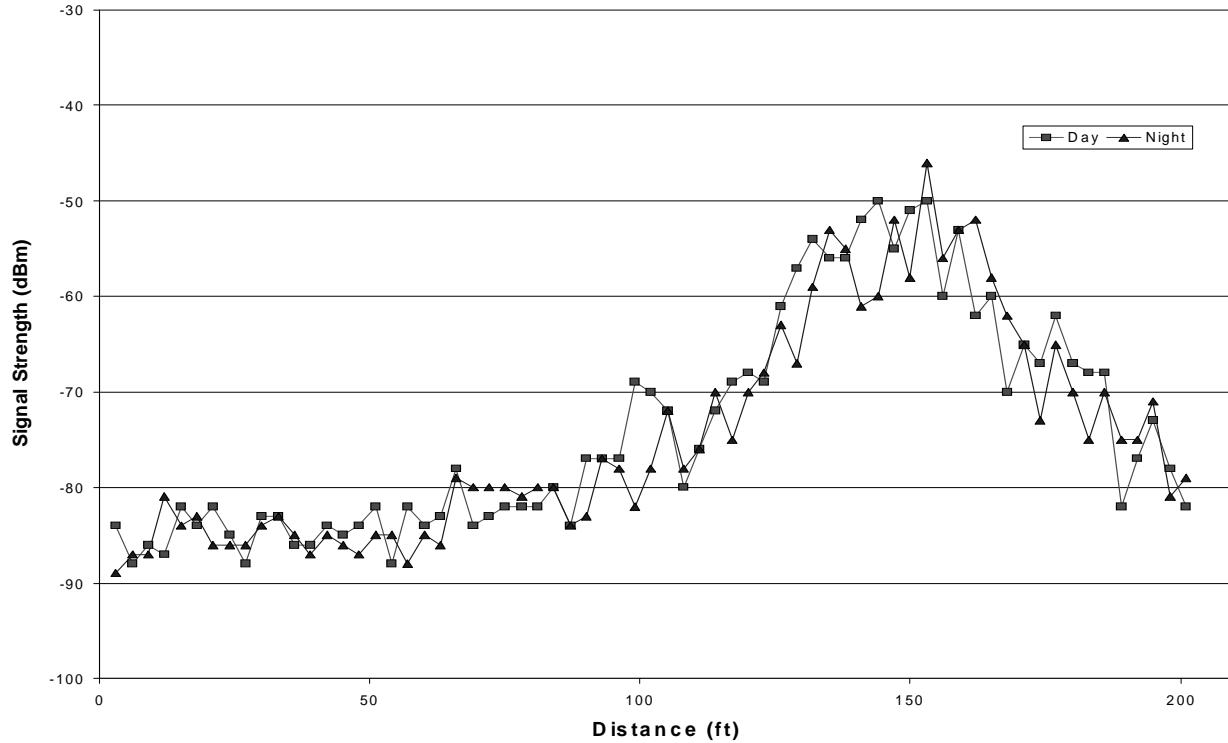


Figure 6. Signal strength comparisons for night and day over a distance of 200 feet

87.5 % of the computed locations were correct, with 97.5 % falling within 12 feet and 100% falling within 26 feet of the correct answer. Our results are better than those presented in [6], with 97.5% of our readings (within 12 feet) as accurate as their 25th percentile (within 11 ft.). The main reason is because we always had line of sight to one access point and averaged more samples when determining location (17 in our case vs. 3 samples, as reported in [6]).

3.2. Table - based: Low Resolution

The low resolution method differs from the others in

% Of Measurements	Error (feet)
87.5	0
90	4
95	5
97.5	12
100	26

Table 2. Accuracy of reported location measurements

that it does not use triangulation. On some platforms, it is not possible to obtain signal strength measurements directly for use. For this reason, a monitoring computer can be used to determine the location of every mobile computer.

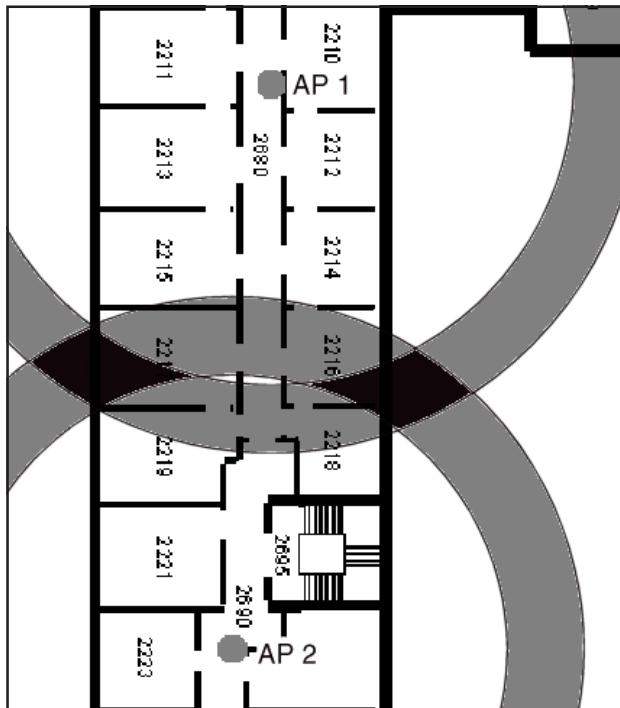


Figure 7. Determining user location

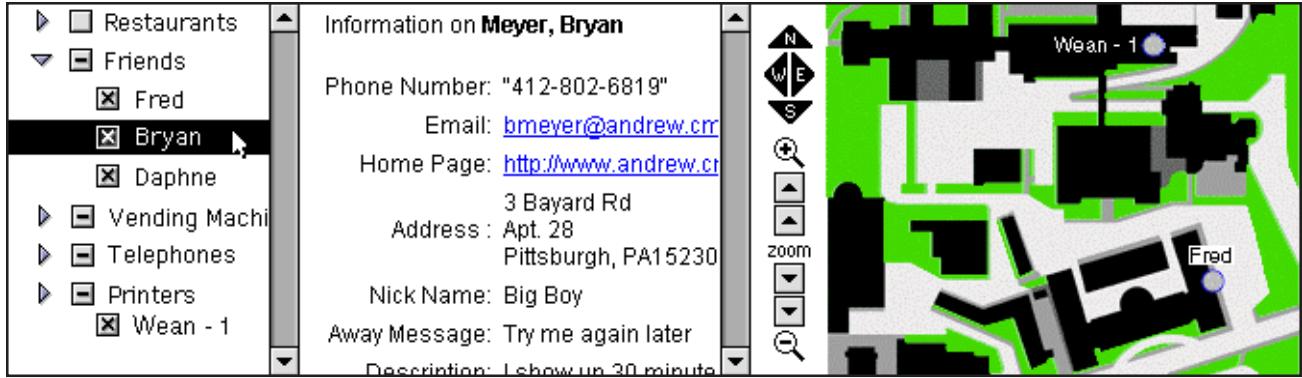


Figure 8. Personal Help Desk application

Each access point acts as a bridge. Using SNMP [7], the access point can be queried for its forwarding table. The table contains a list of all devices that are in contact with that access point. While this method does not provide very high resolution, it can tell us when a mobile device is within the 80 ft. radius covered by each access point. This provides the location of another user without the need to query the user's device directly.

3.3. Model - based

It is also possible to determine location through the use of a signal propagation model. While this is good in theory, it is often too difficult to implement in practice if a high degree of accuracy is required. These models only work for line of sight with nothing in proximity of the measuring computer. In reality, there are obstructions, which attenuate the signal as well as reflecting surfaces that can give a higher signal strength at a location further away from the access point.

3.4. Neural Network - based

An additional method, involving neural networks was explored. The data depicted in Figure 4 was used as two inputs into a neural network with the signal strength from each of the two access points as inputs (26 locations). The network performed very well for over 50% of the locations, with an error of less than one foot when the training data was tested on the network. The other 50% deviated by as much as 20 feet from the optimal answer. We believe this is due to the fact that the training data includes points, which are not strictly increasing or decreasing. Because of this poor performance we did not feel the need to test the network with data other than the training set.

We could solve this problem by customizing the training data so it is strictly increasing or decreasing, but there are times where readings should not be like this. While

neural networks do have the potential to be the fastest method to use, we do not use them because of the errors.

4. Application

Figure 8 illustrates a screendump from an application which has been written using the low resolution method. This application, called Personal Help Desk (PhD) allows a user to determine the location of other users on campus as well as information about them. It also provides other services such as notifying the user of the closest available printer or where food might be available. A server contains a database with information related to users and services and will provide this information upon request.

5. Conclusions and future work

Experiments have shown that on a small scale, the table based method of determining location works very well. The results we have obtained are very promising and will be very useful for future applications.

To cover an area of one acre at one sample every 10 feet, we would require 441 samples. If we were to generate a table to determine a user's location in a large area such as Carnegie Mellon's 103 acres, we would require 44,100 samples. This introduces two problems. The first is the problem of taking these samples. This averages to approximately 110 samples per access point. The second of these problems is searching through all of the samples to determine the location.

The second problem can be addressed by splitting the tables up to cover a smaller area. Because we know which access point is being used for communications, we immediately know the general area a user is in. Each table used can be sized to cover an area containing that access point and all other access points within a 150 foot radius. By allowing the data in different tables to overlap, search time

can be significantly reduced.

In order to solve the problem of taking thousands of samples, we must determine the minimum amount necessary to obtain a good reading. From there we can interpolate the values in between. The minimum set of locations to take values in order to get good results are directly at each access point and at 40 and 80 feet to each side. As stated earlier, all experiments conducted so far have been where there is line of sight, along the length of a hallway. Future experiments will take place in a two dimensional plane and will then extend to three dimensions. Here we would require 20 samples per access point, resulting in 8000 samples total. While this is still a large number, it can be reduced if we do not need as much accuracy.

We are working to increase our resolution. This could be accomplished by taking more samples. While the network card driver only reports signal strength in integers, by taking multiple samples and looking at the distribution of those samples, we should be able to make decimal measurements.

We will also provide a graphical user interface for use with this location tool. It will allow users to see their position, as well as enter correction data when the display is showing an incorrect location.

This paper presents a very comprehensive set of results for using signal strength from wireless access points to obtain location information. This information could be used for obtaining location data as well as other applications.

6. Acknowledgments

We would like to thank Lucent Technologies and Carnegie Mellon Computing Services for their assistance with this project. The research reported in this paper is supported by National Science Foundation, DARPA, and the Institute for Complex Engineered Systems at Carnegie Mellon University.

7. References

- [1] G. D. Abowd, "Software Design Issues for Ubiquitous Computing," Proc. IEEE WVLISI '98, Orlando, FL, April, 1998, pp. 104-109
- [2] B. Brewin, "White House Frees up Civil GPS Signal," <http://www.cnn.com/2000/TECH/computing/05/02/civil.gps.idg/index.html>, May2, 2000
- [3] J. Dorsey, Internal Document, Carnegie Mellon University, May, 2000
- [4] H. Aoki, B. Schiele, A. Pentland, "Realtime Personal Positioning System for a Wearable Computers," Proc. ISWC '99, San Francisco, CA, October, 1999, pp. 37-43
- [5] A. R. Golding, N. Lesh, "Indoor Navigation Using a Diverse Set of Cheap, Wearable Sensors," Proc. ISWC '99, San Francisco, CA, October, 1999, pp. 29-36
- [6] P. Bahl, V. N. Padmanabhan, "RADAR: An In -Building RF-Based User Location and Tracking System," Proc. IEEE Infocom 2000, Tel Aviv, Israel, March, 2000
- [7] E. Decker, P. Langille, A.Rijsinghani,K. McCloghrie, "Definitions of Managed Objects for Bridges," RFC 1493, Network Working Group, <http://www.faqs.org/rfcs/rfc1493.html>, 1993